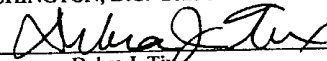


100732808 1071501



PATENT
5181-41200

"EXPRESS MAIL" MAILING
LABEL NUMBER EL424764370US
DATE OF DEPOSIT OCTOBER 8, 1999
I HEREBY CERTIFY THAT THIS
PAPER OR FEE IS BEING
DEPOSITED WITH THE UNITED
STATES POSTAL SERVICE
"EXPRESS MAIL POST OFFICE TO
ADDRESSEE" SERVICE UNDER 37
C.F.R. § 1.10 ON THE DATE
INDICATED ABOVE AND IS
ADDRESSED TO THE ASSISTANT
COMMISSIONER FOR PATENTS,
BOX PATENT APPLICATION,
WASHINGTON, D.C. 20231


Debra J. Tix

Fan Control Module for a System Unit

By:

Jeremy B. Rolls
Michael J. Bushue
Rhod J. Jones
Stepan Tatulian

B. Noel Kivlin
Conley, Rose & Tayon, P.C.
P.O. Box 398
Austin, TX 78701

SUMMARY OF THE INVENTION

Particular and preferred aspects of the invention are set out in the accompanying independent and dependent claims. Combinations of features from the dependent
5 claims may be combined with features of the independent claims as appropriate and not merely as explicitly set out in the claims.

In accordance with one aspect of the invention, there is provided a fan control module for a system unit. The fan control module comprises power outputs for supplying
10 power to a plurality of fan. It also includes a temperature sensor for giving a temperature signal. It further includes a control unit connected to receive the temperature signal. The control unit includes preprogrammed control information for determining power signals to be supplied to each of the fan units for controlling the speed thereof dependent upon the temperature signal.

15 The provision of a separate fan control module for controlling the fan units in a coordinated manner enables reliable and effective cooling of the system unit under widely varying parameters. It also means that existing system components can be employed in harsher temperature environments than they were originally designed for,
20 without needing a complete redesign thereof.

Moreover, where the fan control module includes one or more power inputs from a power supply that is also used to power the other components of the system unit, the fan control module can be provided with electrical noise isolation circuitry to isolate
25 other components of the system unit, from electrical noise generated by the fan units.

In order to limited the power handling requirements of the fan control module circuits, in an embodiment of the invention the fan control module can be logically split into two parts. A first part controls a first pair of fan units and the second part controls a

second pair of fan units. Each part of the fan control module can be provided with respective inputs, outputs and control units. The control information programmed in the control unit of each part can be identical. Preferably, one temperature sensor is employed by both parts to provide a co-ordinated ramp for the fan speeds. Also,
5 where more than four fans are provided, more than two fans per part could be controlled and/or more parts could be employed, as appropriate.

The fan control module is preferably configured on a single circuit board. This provides particular advantages where the fan control card is to be provided as an
10 addition to a system. The temperature sensor is preferably mounted on the circuit board, although it could be placed at some another part of the system as appropriate. Preferably one temperature sensor is used as this facilitates the provision of a controlled and co-ordinated ramp up of the fan speeds. However, more than one temperature sensor could be used, if desired, with each temperature
15 sensor providing respective signals and control of the individual fans being dependent upon individual temperature signals or a function of some or all of the temperature signals.

Preferably speed signals, for supply to an alarms module, are directed via the fan
20 control module and a power distribution board. The fan control module does not process these signals, but the feeding of the signals via the fan control module enables an efficient wiring loom to be made, with a single bundle of wires and a single connector being connected to a fan unit.

25 In accordance with another aspect of the invention, there is provided a system unit including a fan control module, the fan control module comprising power outputs for supplying power to a plurality of fan units, a temperature sensor for giving a temperature signal, and a control unit connected to receive the temperature signal and including preprogrammed control information for determining power signals to be

supplied to each of the fan units for controlling the speed thereof dependent upon the temperature signal.

In a particular embodiment the system unit is a computer system unit including at least one processor module. It may contain anywhere between one and four processor modules. This puts further demands on the cooling requirements, as these will vary in accordance with the number of processors present. Accordingly, the power supply signals output by the control unit can be made dependent upon to the number of processor modules present.

10

In accordance with a further aspect of the invention, there is provided a method of controlling cooling of a system unit, the method comprising:

- a fan control module receiving a temperature signal from a temperature sensor;
- the fan control module determining power outputs to the fan units for
- 15 controlling the speed thereof dependent upon the temperature signal from the temperature sensor and preprogrammed control information for determining power signals to be supplied to each of the fan units for controlling the speed thereof.

In the particular embodiment mentioned above, the system unit is a computer server intended to be rack-mounted for a telecommunications application. It will be appreciated that this puts further strain on the cooling requirements, due to different possible configurations of adjoining equipment in a particular installation, and the possible proximity of other heat generating elements. It will be appreciated that the present invention provides particular and important technical advantages when

25 applied to the adaptation of systems to meet the strict reliability and temperature requirements of, for example, telecommunications applications and that it is ideally suited to such telecommunications applications.

Figures 10A, 10B, 10C and 10D are rear, top, front and perspective views of a power sub-frame for receiving three power supply unit, and Figure 10E illustrates connections for various connectors of a power sub-frame assembly;

5

Figure 11 is a schematic diagram of circuitry from a power distribution board of the power sub-frame of Figure 10;

Figure 12 illustrates the location of an alarm circuit;

10

Figure 13 is a schematic block diagram of the logic of the alarm circuit; and

Figure 14 is a schematic diagram illustrating the configuration of a fan control module.

ALL INFORMATION CONTAINED HEREIN IS UNCLASSIFIED
DATE 03-07-2008 BY 60322 UCBAW

10 more media drives.

In the present instance, two media drives including a digital audio tape (DAT) drive 50 and a CD-ROM drive 52 are provided. Appropriately configured metal cover plates 54 and 56 are provided for the media drives 50 and 52. A disc bay assembly 58 provides a small computer system interface (SCSI) backplane and cables for receiving one or more SCSI media drives, such as a SCSI disc drive 60. Although, in the present instance, the drives are controlled via a SCSI-type interface, it will be appreciated that another media drive interface (e.g., IDE) could be used. A SCSI card (not shown) is located within the chassis to the front of the motherboard. A bezel (decor panel) 62 is provided for covering ventilation holes 63 in the front wall 12 of the chassis 11. A bezel 64 is provided for covering the media drives 50, 52 and 60.

25 A fan control module 66 controls the operation of processor fans 68 and system fans 70. A power sub-assembly that includes a power sub-frame 72 with a power distribution board assembly, is provided for receiving three separate power supply units 74. An alarms module in the form of an alarms card 78 enables the signalling of alarms to the outside world, and is also connected to an LED card 2 for

An isolated ground system is needed in some telco applications when operating in a Regional Bell Operating Company (RBOC) mode. When operating in such a mode, the chassis and logic ground are connected at a remote location to provide, for example, lightning protection. In this case two-hole lugs 101 having a pair of holes 111 to fit over the pair of grounding studs 105 and 107 are provided for each of the power supply units 74 and are secured over the studs using nuts 104 and 106. A similar two-hole lug 101 is secured to the grounding studs 108 and is secured with similar nuts. Earthing wires 109 from each of the two-hole lugs 101 on the power units and the chassis then are taken to the remote, earthing location. The studs 103 105, 107 and 108 are of a standard thread size (M5). The studs 105/107 and the studs 108 are at a standard separation (15.85mm). The studs 105/107 are self-retaining in the insulated board on the power supply units. The stud 103 is self-retaining in the casing of its power supply unit 74. The studs 108 are also self-retaining in the system unit chassis.

In a non-isolated ground situation, chassis ground can simply be tied to a desired ground potential (for example, to the racking system) by connecting a grounding cable to grounding studs 108 provided on the rear of the chassis. A further earth connection is provided via the power cables for the power supplies.

Figure 6 also illustrates rear ventilation holes 110 through which air is vented from the system. Figure 6 also shows the alarms module 78 with a serial connector 112 enabling connection of the alarms module to a network for the communication of

faults and/or for diagnostic operations on the unit 10 to be performed from a remote location. Figure 6 also shows a number of PCI cards 84 received within respective PCI slots 116. A number of further external connections 114 are provided for connection of serial connections, parallel connections and SCSI connections, and for
5 the connection of a keyboard or a Twisted-Pair Ethernet (TPE) connector.

Figure 7 is a plan view of the motherboard 40 shown in Figure 4. Four CPU module slots 120 are provided. Each of these slots is able to receive one processor module 42, and any number between one and four slots may be occupied by a
10 processor module 42. A connector arrangement 122 is provided for receiving a riser card 44 as shown in Figure 4. Also, connectors 124 (in four banks) are provided for receiving DIMMs 46 as mentioned with reference to Figure 4. Edge connectors 126 are provided for connecting the motherboard to connectors mounted on the mounting plane 41. Also shown in Figure 7 is the slot 128 for the alarms
15 module 78 and various ports 130 for the connectors 114 shown in Figure 6.

Figure 8 is a schematic overview of the computer architecture of the system 10. As shown in Figure 8, various components within the system are implemented through application-specific integrated circuits (ASICs). The system is based round a
20 UltraSparc Port Architecture (UPA) bus system that uses a Peripheral Component Interconnect (PCI) protocol for an I/O expansion bus. The CPU modules 40.0, 40.1, 40.2, 40.3, and a UPA-TO-PCI (U2P) ASIC 154 communicate with each other using the UPA protocol. The CPU modules 40 and the U2P ASIC 154 are configured as UPA master-slave devices. An Address Router (AR) ASIC 154
25 routes UPA request packets through the UPA address bus and controls the flow of data to and from memory 150 using a Data Router (DR) ASIC 144 and a switching network 148. The AR ASIC 154 provides system control. It controls the UPA interconnect between the major system components and main memory.

The DR ASIC 144 is a buffered memory crossbar device that acts as a bridge³ between six system unit buses. The six system unit buses include two processor buses, a memory data bus and two I/O buses. The DR ASIC 144 provides crossbar functions, memory port decoupling, burst transfer and First-in-First-Out (FIFO) data read functions. Clock control for the operation of the processor is provided by a Reset, Interrupt, Scan and Clock (RISC) ASIC 152.

The PCI bus is a high performance 32-bit or 64-bit bus with multiplexed address and data lines. The PCI bus provides electrical interconnection between highly integrated peripheral controller components, peripheral add-on devices, and the processor-memory system. A one-slot PCI bus 155 connects to a PCI device 156.0. A three-slot PCI bus connects to three PCI slots 156.1, 156.2 and 156.3. Two controllers are also connected to the second PCI bus 157. These include a SCSI controller 174 and a PCI-T0-EBus/Ethernet controller (PCIO) 158. The SCSI controller 174 provides electrical connection between the motherboard and separate internal and external SCSI buses. The controller also provides for SCSI bus control. The PCIO 158 connects the PCI bus to the EBus. This enables communication between the PCI bus and all miscellaneous I/O functions as well as the connection to slower, on board functions. Thus, the PCIO enables the connection to an Ethernet connection via a Transmit/Receive (Tx/Rx) module 161 and a network device (ND) module 162.

An EBus2 159 provides a connection to various I/O devices and internal components. Super I/O 164 is a commercial off-the-shelf component that contains two serial port controllers for keyboard and mouse, an IEEE 1284 parallel port interface and an IDE disk interface. The super I/O drives the various ports directly with some electromagnetic interference filtering on the keyboard and parallel port signals. The alarms module 78 interfaces with the motherboard and provides various alarm functions. The NVRAM/TOD 168 provides non-volatile read only

memory and the time of day function. Serial port 170 provides a variety of functions. Modern connection to the serial port 170 enables access to the Internet. Synchronous X.25 modems can be used for telecommunications in Europe. An ASCII text window is accessible through the serial port on non-graphics systems.

- 5 Low speed printers, button boxes (for computer aided design applications) and devices that function like a mouse are also accessible through the serial port. The serial port includes a serial port controller, line drivers and line receivers. A one-Mbyte flash programmable read only memory (PROM) 172 provides read only memory for the system.

10

Figure 9 is a perspective rear view of the system 10 showing the withdrawal and/or insertion of a power supply unit 74 in a non-isolated ground situation. In this example, AC power supply units 74 are shown. It can be seen that the power supply unit 74 is provided with the handle 94. As shown in Figure 9, the handle 94

- 15 is provided with a grip 184, a pivot 182 and a latch 180. To insert the power supply unit 74 it is necessary to slide the power supply unit into the power sub-frame 72 with the grip 184 of the handle 94 slightly raised so that the detent 180 can be received under the top 184 of the power sub-frame 72. As the power supply unit 74 reaches the end of its movement into the power sub-frame 72, connectors
- 20 (not shown) provided on the power supply unit 74 make connection with a corresponding connector on the power distribution board at the rear of the power sub-frame 72. Also, at this time, the handle can be pushed down into the position shown in Figure 9. This causes the detent 180 to latch behind the upper portion 184 of the power sub-frame 72. The handle 94 can then be secured in place by
- 25 tightening the screw 95. The AC power supply unit 74 shown in Figure 9 has a single power socket 97, whereas the DC power supply units 74 shown in Figure 6 have two power sockets 96 and 98. Irrespective of whether the arrangement is as shown in Figure 6 with two DC power sockets 96 and 98, or as shown in Figure 9 with one AC power socket 97, the configuration of the power socket(s) and the

lever 94 is such that the lever cannot be moved, and therefore the power supply unit cannot be released from the power sub-frame 72 and the chassis 11 with a plug 186 of a power cable 188 in place in one of the power sockets 96/97/98. The removal operation is achieved by releasing the screw 95, removing the power plug, and
5 lifting and pulling on the handle 94.

In an isolated ground situation, in order to hot-swap a power supply unit 74, it is merely necessary to remove the two-hole lug 101 with its connecting earth wire 109 from the studs 105, 107 of the power supply unit to be removed, to remove the old
10 power supply unit 74, to replace a new power supply unit 74 and then to reconnect the two-hole lug 101 and connecting earth wire 109 to the studs 105, 107 of the new power supply unit 74. These operations can all be performed with the system under power from the other power supply units 74 and with the two-hole lugs 101 and earth wires 109 in place over the chassis studs 108 and the studs 105, 107 of the
15 other power supply units 74.

The isolated ground situation is not shown in Figures 6 and 9. In the non-isolated ground situation shown in Figures 6 and 9, hot-swapping of a power supply unit is even easier, as it is merely necessary to remove the selected power supply unit 74
20 and to replace it with the new power supply unit 74.

Figures 10A, 10B, 10C and 10D are rear, top, front and perspective views of a power sub-frame for receiving three power supply units:

25 The power sub-frame 72 comprises a rectangular, box-shaped frame 191, with four exterior walls on four sides (the top, bottom and two lateral surfaces), one open side 195 for receiving three power supply units and a power distribution circuit board 190 opposite to the open side. In the present instance, the walls are made of electroless nickel-plated mild steel.

Figure 10A shows the power distribution board at the "rear" of the power sub-frame (i.e. opposite to the open side). When inserted in the chassis of the system unit, this "rear" of the power sub-frame is actually the forward-most side of the power sub-frame when viewed with respect to the system unit. The power distribution board 190 is formed with ventilation holes 194 and carries circuit tracks and components (not shown). Figure 10A also illustrates the flanges 198 with screw holes 199 for securing the power sub-frame to the rear chassis wall.

- 10 Figure 10B shows the top of power sub-frame. It will be noted that the power sub-frame body 196 is provided with apertures 197 for lightness and for ventilation purposes.

Figure 10C shows the open (front) side 195 (see Figure 10B) of the power sub-frame. When inserted in the chassis of the system unit, this "front" of the power sub-frame is actually the rear-most side of the power sub-frame when viewed with respect to the system unit. Within the power sub-frame 72, connectors 192A, 192B and 192C for the three power supply units 74A, 74B and 74C, respectively, can be seen. These connectors are mounted on the power distribution board 190 inside the power sub-frame 72. Figure 10C also shows the flanges 198 with screw holes 199 for securing the power sub-frame to the rear chassis wall.

Figure 10D is a perspective view of the power sub-frame 72, which shows that this in fact forms part of a power sub-assembly 71. Internal walls 200 separate three compartments, each for a respective one of the three power supply units 74. Cables 202 connect standby power and signal lines from the power distribution board 190 to a connector 204 for connection to an alarms module. Cables 206 connect main power and signal lines from the power distribution board 190 to various connectors

208, 210, 212 and 214. Figure 10E shows the various connector types 192, 204, 208, 210, 212 and 214 and the electrical signal connections thereto.

Figure 11 is a schematic representation of some of the logic connections on the power distribution board. For ease of explanation, only those connections relevant for an understanding of the present invention are described.

At the left of Figure 11, the three connectors 192A, 192B and 192C for the three power supply units 74A, 74B and 74C are shown. For reasons of clarity and convenience only those connections relevant for an understanding of the present invention as shown. For example, as illustrated with respect to Figure 10E, the connectors 192 have many pins and pass many signals via respective lines. However, as not all of these lines are necessary for an understanding of the present invention, and as it would be confusing to illustrate all of the signal pathways on a diagram, only selected pathways are shown in Figure 11. It is to be noted from Figure 10E, that the power supply units output ground, +3V3, +5V, +12V, -12V and +5V standby potentials as well as control signals such as PSU OK, PSU ON, etc. The +5V standby voltage is used for powering the alarm module 78. The other voltages are for powering the motherboard and other main system components. The various lines could be configured using bus bars, wires, printed circuit or thick film conductors as appropriate.

Firstly, the two-of-three circuit 232 will be explained. This circuit is powered by the +5V standby voltage 231 provided from each of the power supply units 74. Each of the power supply units outputs a PSU OK signal via a pin on its respective connector to a corresponding PSU OK line 230A, 230B and 230C when the power supply unit is operating correctly. Each of these PSU OK lines 230 is connected to the two-of-three circuit 232. This comprises three AND gates 234, 236 and 238,

each for comparing a respective pair of the PSU OK signals. The outputs of the AND gates are supplied to an OR gate 240.

If the output of this OR gate is true, then at least two of the power supply units 74
 5 are operating correctly, and power can be supplied to the motherboard of the computer system. This can be achieved by closing the main power line 245. An output signal 242 could be supplied to a gate 244 (for example a power FET) to enable current to pass to the motherboard and other system components. Additionally, or alternatively, a power OK signal 246 for controlling some other
 10 form of switch mechanism (not shown).

If alternatively the output of the OR gate 242 is false, then this indicates less than two of the power supply units 74 are operative. In this case power is prevented from being passed to the motherboard 40 of the computer system. This can be
 15 achieved by interrupting the main power line 245. An output signal 242 could be supplied to a gate 244 (for example a power FET) to prevent current being passed to the motherboard and other system components. Additionally, or alternatively, a power fault signal 246 could be passed to the alarms module and/or for controlling some other form of switch mechanism (not shown).

20 One-of-three power control is effectively provided by the alarms module 78 to be described later. However, with reference to Figure 11, input A/B signals 268 and output sense signals 270 are passed to the alarms module for standby operation, and control signals 272 could be returned for turning off of a power supply unit, if
 25 required.

Figure 11 further illustrates a protection circuit 256 that is able to detect an overcurrent representative of a current greater than $2 \cdot I_{max}$, where I_{max} is the maximum current that can be output by a power supply, $2 \cdot I_{max}$ being the

maximum current which should be required by the system unit. If a current greater than $2 \cdot I_{max}$ is detected, this is representative of a fault in the system unit. In accordance with telco requirements, in such a situation the system should be powered down. By providing for overcurrent detection on the power distribution board, where the maximum drawable current should be $2 \cdot I_{max}$, it is possible to test for a fault at a lower overall current than if this test were made within each power supply unit. If the test were made in each power supply unit, each power supply unit would need to be tested for an overcurrent in excess of I_{max} , whereby one would be testing for a total current drain of $3 \cdot I_{max}$. This could lead to faults not being detected or not detected early enough and the system could incorrectly be drawing up to $3 \cdot I_{max}$, which could damage components and traces (tracks).

Thus, as shown in Figure 11, each of the main power lines (e.g., +12V) 250A, 250B and 250C from the power supply units 74A, 74B and 74C, respectively is connected to form a common power supply line 254. An overcurrent detector 258 detects a current in excess of $2 \cdot I_{max}$. If such a current is detected (for example as a result of a fault represented by the box 266), then a signal 261 is provided to the connectors 192,A, 192B and 192C for shutting down the power supplies 74A, 74B and 74C, respectively. Also, a signal 262 is passed to a switchable shunt 260 (e.g., a silicon controlled rectifier (SCR), a Metal Oxide Semiconductor Field Effect Transistor (MOSFET), an Insulated Gate Bipolar Transistor (IGBP), etc) to shunt the power supply line 254 to ground. This will cause any energy stored in the power supplies and also in the system (for example as represented by the capacitor 264) to drain to ground, thus protecting the system.

25

The use of the two-of-three circuit described above means that redundant power supply operation is provided in that the system can remain powered even if one of the three power supply units fails. As only two-of-three power supply units are

needed to power the system the third power supply unit can be hot swapped while the other two power supply units power the system.

Figure 11 illustrates the location of an alarms card forming the alarms module 78 in the rear of the system unit 10.

Figure 12 is a functional block diagram for illustrating the alarm sub-system on the alarms module 78. The alarms sub-system provides lights out management or remote management of the system over a serial connection. The alarms module 78 interfaces with the motherboard through an EBus edge connector slot 298 (connected to EBus2 as shown in Figure 8). A PCI-style bracket is attached to one edge of the alarms module (as seen in Figure 11) and provides the external interfaces at the rear of the chassis 11. Internal interfaces provide connections to the power supply assembly and to the LED card 80 located at the front panel of the system unit 10. The alarms module is powered by the standby, or reserve, power supply. The alarms module only requires power from a single power supply to remain operable. Accordingly, the alarms module can remain operable even in a situation where the system has been powered down due to there being only one power supply unit operable.

20

The alarms sub-system comprises a logic device 280 which receives inputs 298 from the EBus, inputs 286 from the fans, input 290 from general purpose events, input 270 from the power supply unit output rails and inputs 268 from the A and B power inlets. The logic circuit samples, or multiplexes, the inputs to the microcontroller 296 in response to multiplex signals from the microcontroller 296. The microcontroller 296 processes the sampled (multiplexed) inputs. The microcontroller 296 provides power control signals 272 for controlling the power supply units, and alarm outputs for the output of alarm signals. The microcontroller 296 also outputs power supply unit status signals 304 and fault signals 306. The

all other things being equal, the more the

Figure 14 illustrates the configuration of the fan control module 66 shown in Figure 4. The fan control module is subdivided into two halves 66A and 66B. One half 66A handles one processor fan 68A and one system fan 70B and the other half 66B handles the other processor fan 68B and the other system fan 70B. The fans are connected to the fan control module 66 by respective power lines 320 so that the fans receive their power via the fan control module. The fan control module receives +12V power via power lines 324A/B from the power distribution board 190 and supplies voltages to the fans via the power lines 320 in a controlled manner.

Sub
B1

sub
BI
conc.

instead forwarded via tacho sense 326 to the power distribution board 190. The power distribution board then routes the tacho sense signals to the alarms module 78 to form the signals 286 shown in Figure 13. This routing is convenient as it enables simpler wiring looms to be used. Also, when replacing a fan unit, the maintenance engineer only needs to remove a single bundle of wires from the fan to the fan control module 66, rather than having to locate a number of different connectors connected to the fan. The fan control module thus has four fan connectors, each for receiving a connector connected to a bundle of wiring from a respective fan, plus a further connector for receiving a connector with a bundle of wires from the power distribution board.

As shown in Figure 14, each half 66A/66B of the fan control module receives respective power lines 324A/B from the power distribution board. Each half of the fan control module includes electrical noise isolation circuitry 340A/B. This electrical noise isolation circuitry 325 A/B, which can be of conventional construction, prevents dirty power signals on the lines 320A/B caused by electrical noise from the fans being passed back along the power lines 324A/B and potentially contaminating the otherwise clean power supply to the electronics of the system unit (e.g., the components on the SCSI bus. The provision of clean power supply signals in a telco application is important in order to ensure reliability of operation. Although in the present example the noise isolation circuitry is located in the fan control module, it could be located elsewhere as long as it is effective to isolate the main power lines from fan-related electrical noise.

As further shown in Figure 14, each side 66A/B of the fan control module comprises control logic 342A/B which receives a temperature signal from a temperature sensor 344 and adjusts the speed of the fans by adjusting the voltage supplied thereto in accordance with pre-programmed parameters in order to provide a desired degree of cooling. The control logic 342A/B can be implemented by an

ALL INFORMATION CONTAINED HEREIN IS UNCLASSIFIED

[REDACTED]

DYC REF: P007764US